TechTrends
https://doi.org/10.1007/s11528-019-00466-3

ORIGINAL PAPER

Spam and Educators' Twitter Use: Methodological Challenges and Considerations

Jeffrey P. Carpenter[1] · K. Bret Staudt Willet[2] · Matthew J. Koehler[2] · Spencer P. Greenhalgh[3]

Abstract
Twitter and other social media have assumed important places in many educators' professional lives by hosting spaces where new kinds of collegial interactions can occur. However, such spaces can also attract unwelcome Twitter traffic that complicates researchers' attempts to explore and understand educators' professional social media experiences. In this article, we define various kinds of spam that we have identified in our research on educators' uses of Twitter. After providing an overview of the concept of spam, we evaluate the advantages and disadvantages of different approaches to addressing the presence of spam in educator-focused Twitter spaces. Then we suggest practical, holistic metrics that can be employed to help identify spam. Through secondary analyses of our past research, we describe the use of such metrics to identify and deal with spam in three specific cases. Finally, we discuss implications of spam and these suggested methods for teacher educators, instructional designers and educational technology researchers.

Keywords Hashtags · Professional community · Professional learning · Research methods · Spam · Social media · Twitter

Researchers are increasingly taking interest in how educators use social media platforms for professional learning purposes. Technology has long been recognized as offering new learning opportunities to educators (e.g., Dede 2006; Macià and García 2016; Tour 2017), and the specific role of social media platforms has recently received increased attention (e.g., Prestridge et al. 2019). Although educators use a variety of social media to advance their own learning, Twitter has been particularly popular among both educators organizing their own professional learning and the researchers studying them (e.g., Carpenter and Krutka 2014; Rosenberg, Greenhalgh et al. 2016; Tucker 2019; Veletsianos 2017).

One of the features of Twitter that makes this platform especially helpful for creating professional learning communities is the *hashtag* (i.e., a word or phrase, preceded by a hash or pound sign, that organizes content by topics). Education-focused Twitter hashtags provide educators with spaces to connect, share ideas and discuss education topics (Carpenter et al. 2018; Rosenberg et al. 2016). These spaces have been evaluated in terms of *communities of practice* (e.g., Britt and Paulus 2016; Gao and Li 2017), *professional learning communities* (e.g., Goodyear et al. 2019) and *affinity spaces* (e.g., Staudt Willet 2019; Rosenberg et al. 2016). Although there are important distinctions between these frameworks, the literature generally concludes that professional learning spaces created through Twitter hashtags can be of real value to educators. It is therefore unsurprising that some teacher educators now explicitly introduce Twitter and Twitter hashtags to preservice teachers (e.g., Carpenter 2015; Carpenter and Morrison 2018; Hsieh 2017; Luo et al. 2017).

Yet, the same Twitter features that allow for the easy creation of spaces for professional learning also allow *spam*—"unsolicited, repeated actions that negatively impact other people"

✉ Jeffrey P. Carpenter
jcarpenter13@elon.edu

K. Bret Staudt Willet
staudtwi@msu.edu

Matthew J. Koehler
mkoehler@msu.edu

Spencer P. Greenhalgh
spencer.greenhalgh@uky.edu

1 Elon University, 336 278 5969, Campus Box 2105, Elon, NC 27278, USA

2 Michigan State University, 620 Farm Ln, East Lansing, MI 48824, USA

3 University of Kentucky, 320 Lucille Little Library Bldg., Lexington, KY 40506, USA

Published online: 04 December 2019

(Twitter n.d.)—to affect educators' online activities. Brunton (2013) has described spam in opposition to community; that is, by affording the concentration of attention in a community, technologies also create opportunities for those who wish to profit in some way from that concentrated attention. For instance, hashtags can serve as simple *portals* (Rosenberg et al. 2016) to particular professional learning spaces in that participants need only search for hashtags to find communities or include hashtags in their tweets to contribute to them. However, this low threshold allows for irrelevant, off-topic, commercial, or malicious messages to easily enter the space. Because popular Twitter hashtags attract attention from a considerable audience, such hashtags can prove an enticing target for spammers who attempt to redirect users' attention. Yet, the topic of spam has so far received little attention from education scholars.

The purpose of this paper is to explore the implications of spam—especially in the form of extreme participation inequality—for research on education-focused hashtags. Researchers who investigate educators' Twitter use have shown great interest in defining, describing and distinguishing the spaces formed through particular hashtags. However, to carry out these objectives, researchers must determine how to address the common presence of spam content. On one hand, keeping spam in a dataset for analysis risks distorting researchers' view of the composition of and interactions within a given space. On the other, excluding spam from an analysis risks misrepresenting the actual experiences of teacher professional learning and may inadvertently contribute to an unwarranted utopian rhetoric around social media (see boyd 2014). Finally, researchers may have an interest in analyzing spam or spam-like behavior. In any of these cases, however, identifying spam and establishing a plan for what to do with it is a key step, albeit an often-overlooked one.

Background

In the following sections, we describe some basic characteristics of spam as well as how personal and contextual factors may influence an individual's perceptions of those characteristics. In doing so, we draw on selected sources from the broader literature on spam while concentrating on how spam and spam-like activity have been acknowledged in Twitter research by educational technology scholars. We conclude with a brief summary of how researchers have responded to spam.

Characteristics of Spam

Spam has been described as "undesirable text, whether repetitive, excessive, or interfering" (Brunton 2013, p. xxii). This broad definition is helpful for acknowledging the wide range of forms that spam can take, but attention to some of the specific ways text can be undesirable (especially within an education-focused community) is necessary for fully understanding this phenomenon. These are described below and exemplified in the faux tweets in Fig. 1.

- *Unintentionality*: The same hashtag may mean different things to different people, leading to two different groups accidentally occupying the same space. Similarly, mistakes or typos may lead to someone posting to an educational hashtag entirely by accident. For example, in Fig. 1a, the "#maet" hashtag is used by participants in a Master of Arts in Educational Technology program, but the functionally-identical "#mæt" hashtag refers to the Danish word for "full" (see Greenhalgh et al. 2016).
- *Irrelevance:* Hashtags are open spaces, which lowers barriers to participation for both educators and those speaking on other, unrelated topics.
- *Commercial nature*: Among the unrelated messages that may enter an educational hashtag are those promoting commercial products (whether educational or not).
- *Offensiveness*: The attention gathered by a Twitter hashtag serves as a tempting target for those seeking to spread offensive material, whether for commercial reasons (e.g., in the case of pornography), shock value (e.g., in the case of violent images or profanity), the hijacking of an unwitting audience (e.g., when attacking or insulting individuals or an entire community) or some other motivation.
- *Impersonality*: The effect of any of the characteristics described above can be made even more undesirable if content is posted through automated or inauthentic means (e.g., *bots*; Mowbray 2014), especially as they allow for more posting of spam content with fewer consequences.

Although it is helpful to consider these characteristics separately, Fig. 1e (a tweet that is both impersonal and irrelevant) demonstrates that they can—and often do—co-exist within a single message.

Participant Responses to Spam

We expect that spam affects educators' participation in Twitter hashtags, but those effects can depend largely on personal and contextual factors. Donath (2007) noted that spam can affect users as they forge social ties and identify legitimate content; however, if spam can be defined as "undesirable text" (Brunton 2013, p. xxii), much depends on how individuals and communities define "undesirable." A community may all be on the same page in their judgment; for example, we have observed one case (Rosenberg et al. 2017) in which pornographic spam led an entire community to temporarily abandon a hashtag in favor of another (see also Carpenter and Harvey 2019 for a discussion of switching platforms in response to spam). Conversely, in cases where Twitter

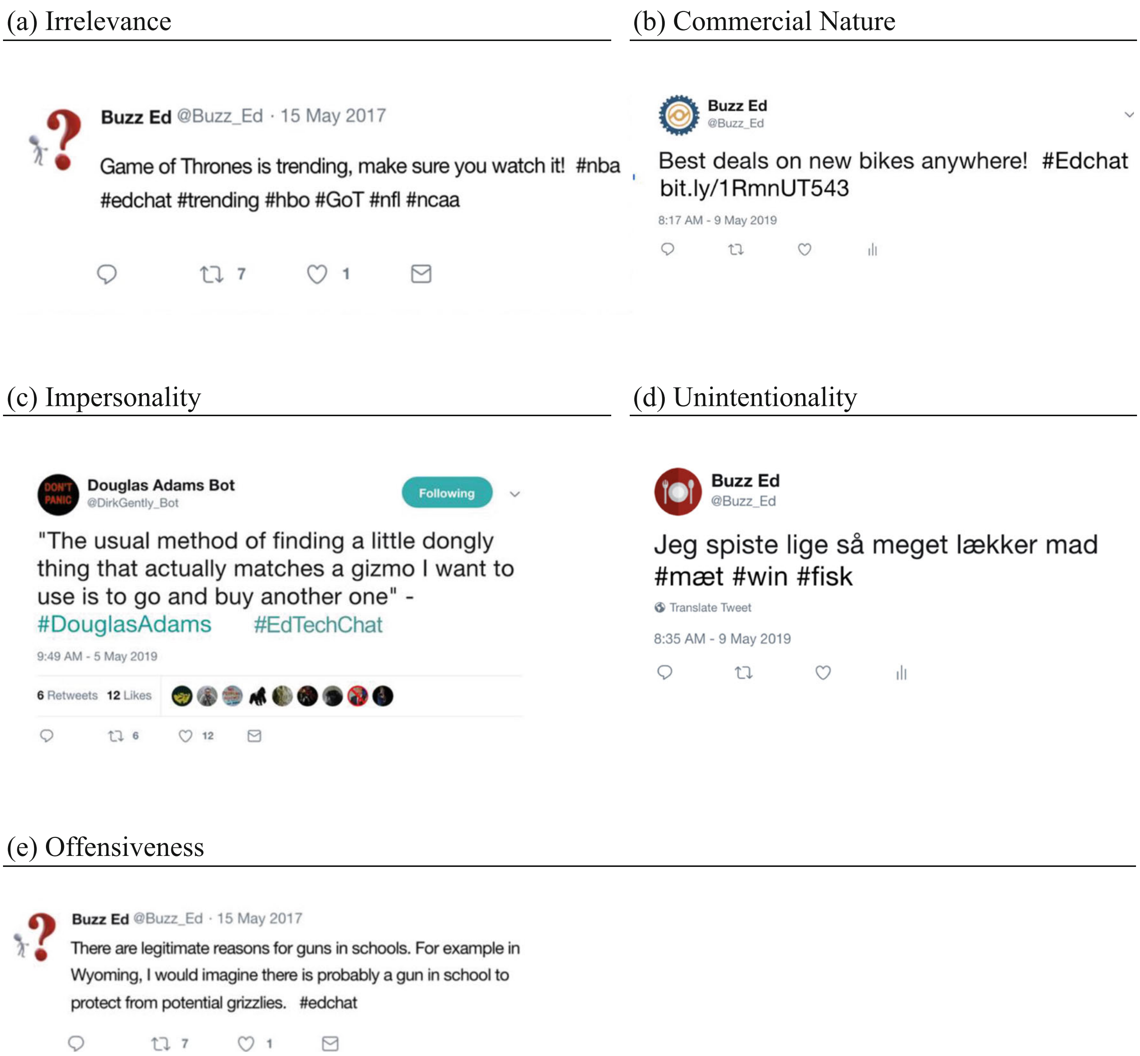

Fig. 1 Artificial tweets as examples of characteristics of spam

communities have an explicit focus on certain products (e.g., #TLAP for Burgess's 2012 book, *Teach Like A Pirate*, or #gafe for Google Apps for Education), what would typically be considered commercial spam may, in fact, be welcomed.

In other scenarios, however, individual cases may be evaluated differently by individual participants. A tweet using several hashtags to promote an educational workshop (see Greenhalgh 2018) may be of genuine interest to some teachers but seen by others as an unwelcome commercial intrusion. Indeed, a tweet could feature none of the characteristics described above but still be deemed "undesirable." For example, many educational hashtags host regular, synchronous *chats*: hour-long conversations structured around pre-determined prompts (Evans 2015; Gao and Li 2017). Research has reported that some participants feel overwhelmed during high-volume chats (Britt and Paulus 2016; Luo et al. 2017); even though this may not correspond with one's intuitive understanding of spam, the sheer volume may still discourage participation just as offensive or commercial content would.

Researcher Responses to Spam

Researchers may also differ in their perceptions of spam—more importantly, they face the task of consistently identifying and intentionally responding to spam. As the bulk of this paper is dedicated to our suggestions for identifying spam, we focus here on the different ways that researchers may respond to spam. In some instances, spam itself may even be of interest. For instance, Twitter has been plagued by cyberviolence and misogyny (Nagle 2018), and researchers may be interested in documenting the degree to which such *offensive* or harassing spam is present in ostensibly education-focused Twitter spaces. Furthermore, the phenomenon of *teacherpreneurship* (e.g., Shelton and Archambault 2018) may lead researchers to investigate commercial messages originating with teachers. However, spam also raises the uncomfortable possibility that a researcher's "findings are an artifact of a flawed dataset" (Karpf 2012, p. 642). Researchers seeking to describe "normal" behavior may judge that spam is not only undesirable for participants but also unwanted "noise" in their data (Kwak et al. 2010). For example, Greenhalgh (2018) acknowledged that his description of one specific hashtag was skewed by the unintentional participation of an individual with a large network.

A Practical, Holistic Approach to Spam in Educational Twitter Research

In this section, we discuss scenarios in which education researchers may encounter, identify and consider how to treat Twitter spam. Previously, the science of spam detection has guided research towards automated approaches that seek to identify either individual spammers or individual spam tweets (Wang et al. 2015). These computational approaches aim to solve a *classification* problem—the goal is to create an algorithm that correctly sorts users or tweets as spam or not spam. These are often resource-intensive approaches, requiring access to specialized training in machine-learning techniques, large training datasets of previously classified tweets and computational resources (e.g., Lin and Huang 2013; Wang et al. 2015). The use of machine-learning approaches for addressing spam in educational social media research can be associated with two problems. First, spammers and spam-detection algorithms are constantly evolving—often, the newest detection approaches rapidly decrease in effectiveness as spammers devise new strategies to circumvent the latest detection algorithms (Chen et al. 2016). Second, machine-learning algorithm approaches are not currently readily available to the many educational researchers who lack access to and training in the latest computer science methods (Kimmons and Veletsianos 2018).

In contrast to such complex and resource-heavy machine-learning detection methods, we argue here for an approach that focuses on practicality for educational researchers. We first consider a number of metrics—that is, quantitative measures that may help identify possible spam users at scale. Then, we focus on additional considerations that should be weighed alongside these metrics for more holistic decision-making.

Practical Metrics for Educational Researchers

The practical approach described herein relies upon several metrics that are straightforward and readily available to educational researchers working with Twitter data. These metrics are not themselves direct indicators of spam but can be utilized in combination to identify potential spam for further scrutiny. This set of metrics should not be considered exhaustive; instead, it offers a starting point for a discussion of the practicalities of dealing with spam. These metrics primarily identify spam at the level of the user.

- *Volume of tweeting*: One spam indicator is unusually high-volume tweeting, which is often bot generated. This could be measured by a raw count of tweets, or the percentage of tweets posted to a hashtag by a single user.
- *Level of interaction*: Because spammers often broadcast messages that others ignore (Lin and Huang 2013), spam accounts can also be identified by the absence of other users' interaction with their tweets. Researchers can therefore examine the extent to which a suspected spammer's tweets result in likes, retweets and replies.
- *Following* vs. *followers*: Spammers frequently have imbalanced following/follower ratios, as they follow large numbers of other users but themselves have comparatively few followers.
- *Level of hashtagging*: Some spammers include numerous hashtags in their tweets to promote their content to as broad an audience as possible. However, it is not uncommon for regular, non-spammer tweeters to also include more than one hashtag, so this metric should be considered in light of the norms of the hashtag space being studied. Researchers could consider raw number of hashtags, the percentage of tweets that contain more than one hashtag (i.e., a hashtag in addition to the one defining the space under investigation), or the average number of hashtags per tweet.
- *Level of hyperlinking*: Spammers often include hyperlinks in their tweets in an attempt to drive traffic to certain websites (e.g., Lin and Huang 2013). For instance, a tweet may advertise goods for sale and include a hyperlink to the website where those goods could be purchased. Researchers can therefore analyze the raw number of links, the average number of links per tweet, or the percentage of tweets that include links.
- *Bot-like activity*: Automated accounts (i.e., bots) can be the sources of a lot of the commercial spam found in educator social media spaces. Researchers and Twitter, Inc. continuously try to develop tools and techniques to detect bots (Chen et al. 2016), while many bot creators simultaneously change techniques to thwart detection. Tools such as *Botometer* (https://botometer.iuni.iu.edu/) examine a user's Twitter activity and estimate the probability that the user is a bot.

These metrics do not align directly with the types and characteristics of spam that we outlined earlier. Characteristics such as *offensiveness* and *irrelevance* are often judgement calls that require interpretation and context not easily detectable in large Twitter datasets. Furthermore, as previously described, personal and contextual factors play a key role in the difference between potential spam and actual spam. Nonetheless, *irrelevant, commercial, offensive,* and *impersonal* forms of spam frequently leave behind digital traces that are reflected in the above metrics. Thus, these metrics can serve as a starting point that can flag potential spam; the next section describes some of the considerations needed for a more holistic evaluation.

Additional Considerations for Holistic Evaluation

We suggest identifying spam by using human raters who look across the multiple metrics we have described above to arrive at holistic decisions regarding whether particular data should be included in a study's dataset. This contrasts with computational methods of spam-detection, in which machine-learning algorithms rely uniquely upon metrics—some of which are similar to the ones mentioned above—to arrive at decisions regarding whether content is spam. Given the diversity of ways in which Twitter is used by educators, exclusive reliance upon one quantitative metric (or even several) would likely result in failures to identify some spam while also falsely identifying some accounts as spammers.

A holistic identification of spam may involve the use of qualitative tests that allow for more nuance than quantitative metrics. For example, simple review of Twitter profiles and a few recent tweets can often serve to triangulate the aforementioned metrics, as these data can indicate if accounts belong to educators or at least seem linked to education-related purposes. In one recent study (Carpenter et al. 2018), a single user accounted for a suspiciously high percentage of the traffic for #TLAP, a hashtag associated with the book *Teach Like a Pirate* (Burgess 2012). However, checking that user's profile revealed that the account belonged to the book's author, and it was thus logical that he was tweeting frequently with the hashtag.

Three Examples of this Practical, Holistic Approach in Use

Although the approach described in the previous section is useful for identifying spam, it does not guide researchers' decisions in what to do after having identified spam. These decisions should align with the research questions under investigation, the nature of the data itself and the level of certainty in the decision being made. For instance, research questions related to the full range of experiences, uses and purposes of an education-focused hashtag are likely to exclude very little data from analysis. After all, some spam can unfortunately be part of educators' typical Twitter experiences. However, if research questions are directed towards the professional interactions that Twitter can facilitate, spammers who fail to actually engage with educators could reasonably be excluded from the dataset.

In a previous conference paper (Carpenter et al. 2019), we described the use of this holistic approach to identify spam in one dataset. In this section, we describe different practical benefits of identifying this spam and the associated considerations researchers can make. In doing so, we expand our description of that instance of spam identification and introduce two other implementations of our proposed approach. We present these examples of identifying spam and making decisions about its inclusion not to suggest a one-size-fits-all approach for educational researchers, but instead as an opportunity for scholars to consider appropriate metrics and strategies related to the spam that they may encounter in their own research. A holistic decision-making process can lead researchers in different directions, depending on a study's focus and research questions.

Practical Use #1: Allowing for Accurate Comparisons

In some studies, being able to accurately identify and remove spam can be essential for making valid comparisons between communities. Carpenter and colleagues (2018) aimed to depict the landscape of education-focused Twitter hashtags by comparing and contrasting the traffic from a set of 16 hashtags over a 13-month timespan. However, the nature and amount of spam associated with the different hashtags made comparing the hashtags more difficult. Some of the 16 hashtags were characterized by relatively little undesirable content, but multiple hashtags appeared to feature large quantities of spam. For instance, for #BFC530, one account posted more than 30,000 tweets during one three-month time period. Such outliers challenged attempts to make quantitative comparisons of the hashtag traffic that represented actual educators' use of those hashtags.

In response, Carpenter and colleagues (2018) conducted a multi-stage spam removal process that resulted in 28 accounts being scrutinized to determine if they should be removed from the dataset. Using the metrics described previously, the research team holistically analyzed each of these users to determine if their tweets should be excluded or not. In some cases, the decision was relatively straightforward. For example, one account sent 7414 tweets including a particular hashtag, every one of which contained a link and none of which had been retweeted. This prevalence of linking appeared to be bot-like behavior, and the absence of any retweets suggested that the account's content was not considered relevant or valuable by other users.

The metrics were not, however, always conclusive. In some cases, it was only by the researchers' manually reviewing tweets sent from these accounts that a decision could be made. For example, one individual sent a suspiciously large volume of tweets to the #satchat hashtag; however, the percentage of those tweets that included hyperlinks was not abnormal, the tweets had received over 1000 retweets and a bot-identification tool suggested there was a very low probability the account was a bot. Nonetheless, a review of recent tweets sent from the account revealed that tweet content was focused almost entirely on politics, much of which could be considered conspiracy theories (e.g., Pizzagate), and little of which related to education. It seemed likely, therefore, that the user was simply attracted to #satchat because of the attention already concentrated there and that their off-topic tweets were largely retweeted by fellow conspiracy theorists, rather than by educators.

These spammers had important effects on the study's research questions. For example, Carpenter et al. (2018) sought to compare trends in these hashtags over time. Another account active in the #satchat hashtag appeared to be a bot and during the last 3 months composed more than 45,000 #satchat tweets. This flurry of spamming meant that prior to spam removal, preliminary analysis suggested that #satchat traffic had increased by 107% between September 2016 (28,247 tweets) and September 2017 (58,451 tweets), the largest percentage growth among the 16 hashtags. However, once spam had been removed, the increase was instead 44%, from 28,247 tweets to 40,705 tweets (Fig. 2), and four other hashtags had experienced greater percentage growth.

Similarly, preliminary analysis (i.e., prior to spam removal) suggested that the #BFC530 hashtag featured dramatically different usage patterns in comparison to the other hashtags in the sample. The other 15 hashtags averaged between 3 and 10 original tweets per month per user, but it appeared that #BFC530 averaged approximately 26 original tweets per month per user. This seemed to suggest that #BFC530 featured an intensity of posting behavior that was quite unique in comparison to the rest of the hashtag sample. However, we identified that nine of the top 10 #BFC530 users were spammers, and the removal of those spammers and their tweets meant that the original tweets per user number for #BFC530 decreased to approximately 8.5 per month, within the range of other hashtags studied. The removal of spam therefore enabled comparisons of the 16 hashtags that were more valid given the research questions being investigated.

Practical Use #2: Adding Confidence to Findings

A holistic approach to identifying spam can add confidence to research findings even if that spam is ultimately not removed. Participation inequality among users of social media is a commonly reported phenomenon, with the great majority of activity coming from a few participants—even to the extent of contributions following a power law distribution (Gao and Li 2017; Kraut and Resnick 2011). Because users contribute at such different rates, a few people can affect the overall appearance of the group—this was evident in the case of participation in #satchat and #BFC530, which we previously described. However, not all of the problem cases flagged by quantitative metrics will be determined by a more holistic evaluation to be spam, and not all cases of spam will have a significant effect on a study's findings.

To demonstrate what we mean, we consider the case of another Twitter hashtag—#Edchat. We applied some of the procedures and measures from Carpenter and colleagues (2018) to data from another study (Staudt Willet 2019), starting with 1,228,506 unique #Edchat tweets from 196,263 different tweeters across 8 months. Because prolific tweeters were often found to be spam by Carpenter and colleagues

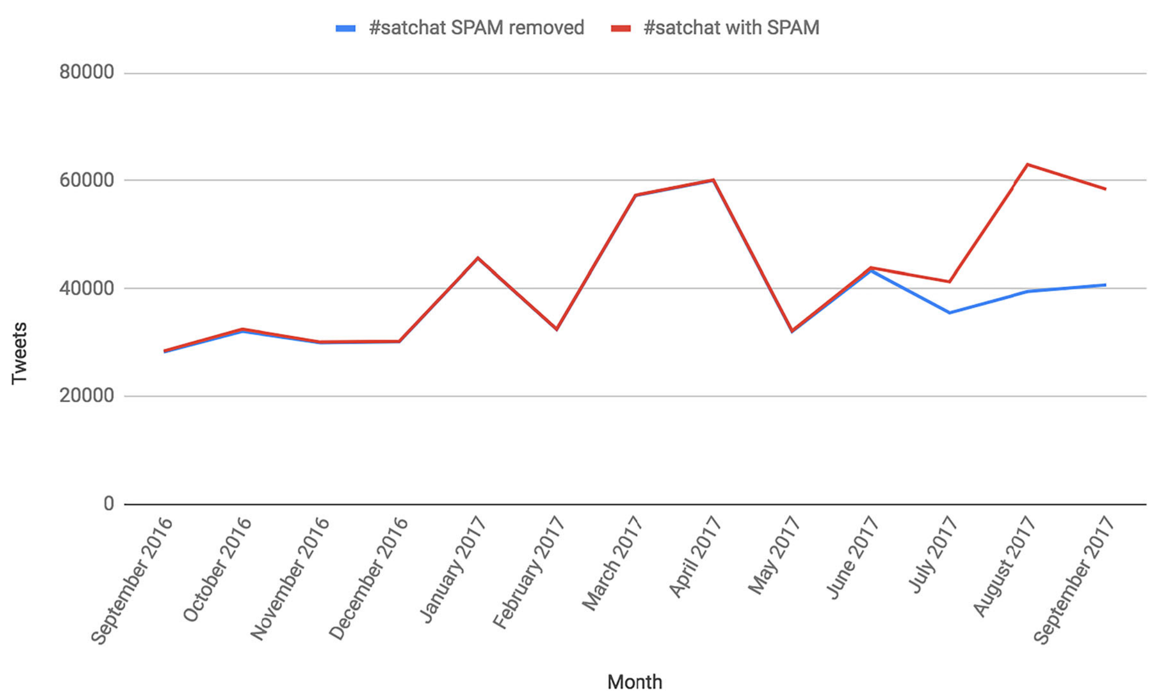

Fig. 2 Tweets per month including #satchat, with and without spam users

(2018), we looked at the most prolific #Edchat tweeters—specifically, the top 10 in terms of tweet volume. The top 10 #Edchat tweeters contributed a very high volume of tweets: the most extreme case was more than 277 standard deviations above the mean, and even the tenth highest tweeter contributed at nearly 50 standard deviations above the mean. Despite these metric-based red flags, a qualitative analysis of profiles and recent tweets demonstrated that these prolific tweeters were generally contributing content relevant to #Edchat—even tweeters whose profiles clearly identified them as bots. In this case, the relevance of tweet content superseded tweet volume or bot status—making the important point that bots should not be automatically considered spammers in all cases.

Even if we had removed these 10 accounts as spammers, there may not have been a difference in findings to the same extent there was in the Carpenter et al. (2018) study. For example, we have previously described how removing spam from the #BFC530 hashtag resulted in a large impact on the original-tweets-per-month-per-user measure. In the case of #Edchat, the 68,905 tweets associated with the 10 prolific #Edchat tweeters accounted for 6.02% of the tweet corpus, and even if all of them had been identified as spam, their removal would have had a less dramatic minimal impact on the original tweets-per-month-per-user measure used by Carpenter et al. (2018): 2.31 with all tweeters included versus 2.10 after hypothetical removal of the top ten users. The exclusion or inclusion of these prolific tweeters in analysis did not as dramatically change the overall measure in the same way as it did in the case of #BFC530 described in the Practical Use #1 section.

Although ultimately we decided not to remove any spammers from the #Edchat dataset, this informed decision could not have been made without first applying the holistic evaluation process. Similarly, in this particular case, a holistic evaluation helped increase confidence that key measures would not be unduly swayed by a few users' activity. We recommend that researchers plan to implement spam detection at the outset of data analysis. In this way, spam metrics can serve as checks for participation outliers and determine how much the prolific contributors affect the overall description of activity (whether or not they are determined to be spammers). Sometimes, as in the cases of #satchat or #BFC530, participation inequality greatly impacts findings; in other cases, such as for #Edchat, the contributions of the 10 most prolific contributors were on topic and did not dramatically impact the description of the overall hashtag traffic.

Practical Use #3: Focusing on Spam-like Behavior

Unlike the previous two cases, we describe here how our practical metrics can identify potential spammers not in order to remove them but rather to focus on them as an actual phenomenon. Throughout these cases, we have focused on how these spam metrics can be helpful in identifying users whose Twitter activity is far outside the norm. Regardless of whether this extreme activity is ultimately considered to be spam, researchers may find value in studying these outliers. For example, Staudt Willet (2019) suggested a need for future research to study self-promotional behavior on Twitter, which may appear to some to be spam. We therefore revisit the #Edchat data described above with the goal of identifying self-promoting users.

Given this goal, our focus is not just on *volume* metrics but also on *interaction* metrics such as the average likes, retweets and replies per tweet. We conceptualized a user-archetype of high volume tweeters (i.e., more than two standard deviations above the population mean) who contribute at least some original tweets (i.e., not just solely retweeting), follow more than twice as many users as follow them in return and whose contributed content receives very little engagement (i.e., their tweets average less than one like, one retweet and one reply). In other words, this archetype represents tweeters who seem to be trying very hard to get noticed but do not appear very successful—the potential profile of someone engaging in excessive self-promotional activity. We found 29 #Edchat tweeters who fit this archetype and whose Twitter behavior could subsequently be studied in further research.

In addition, we considered one specific type of behavior that would seem to be associated with self-promotion: the inclusion of numerous hashtags in each tweet. We considered the mean hashtag-per-tweet score for users who contributed at least some original content, and we found 7897 tweeters who averaged four or more hashtags per tweet, 1675 who averaged at least seven and 568 who averaged more than 10 hashtags per tweet. Other researchers could decide which of these warrant further investigation; our point here is to illustrate the usefulness of our practical metrics for more than just detecting spammers for removal—these metrics can also help identify specific cases for further study.

Discussion

Researchers will have to make different choices about what to do with spam in education-focused Twitter hashtag traffic according to their data and research questions. Different elements of the practical, holistic approach we have outlined may be relevant for particular studies and cases. With time, computational spam detection methods may eventually improve in effectiveness and/or may become more readily accessible to educational researchers. Social media such as Twitter may also improve their own policing of certain types of spam. In the near future, however, we see value for many research contexts in the utilization of a combination of metrics and a final holistic human decision to find spam and consider its removal. Users who access education-focused Twitter hashtags have diverse purposes and motivations (e.g., Carpenter and Krutka 2014), and the multiple flavors of spam that can affect those spaces can confound purely algorithmic methods.

Although tweets to an education-focused hashtag that advertise counterfeit handbags or pornography are relatively easily characterized as spam, spam is not always so self-evident. A fundamental challenge that complicates decisions on whether to remove from a dataset particular tweets or users is the subjective element in defining spam. While some users may feel that tweets that advertise education products detract from Twitter's usefulness, others could see value in the advertised products or could simply accept such tweets as a minor and inevitable nuisance that can be expected when utilizing a non-subscription-based commercial platform. For example, many radio and TV users tolerate advertisements, even if they do not particularly appreciate the advertisements' presence. Spam's consequences may therefore differ depending on users' varied perspectives. Some users may perceive blocking or muting a few prominent spammers to be merely a minor nuisance, and this may be all that is required to limit spam's impact on their experiences with some hashtags. Meanwhile, other educators could find spam to be unsettling and even toxic (Nagle 2018) and, as a result, avoid hashtags that feature large amounts or certain kinds of spam.

Online spaces and communities that attract more participants can offer benefits to educators, including access to more resources, more perspectives and more potential new colleagues (Britt and Paulus 2016); however, bigger is not always better (Carpenter and Morrison 2018; Staudt Willet 2019), particularly given associated spam-related concerns. Spam can be considered as the inevitable abuse of the same technologies that facilitate the formation of such large communities (Brunton 2013). It could prove discouraging for some educators if their first Twitter experiences are solely with the most popular hashtags where they could have to navigate relatively more spam. Additionally, some Twitter spaces feature disorienting quantities of content (Britt and Paulus 2016), even if individual messages do not bear the hallmarks of spam. It could be that Twitter hashtags with smaller volumes of traffic and fewer users are more beneficial spaces because they feature relatively less spam.

The examples of spam removal described above deal with identification at the user level, but identifying spam at the tweet level may be more appropriate or helpful for some studies. For instance, some individuals engage in *teacherpreneurship* via social media such as Twitter (Shelton and Archambault 2018). Such teachers may at times tweet content related to professional learning, but also sometimes advertise the materials they are selling on lesson marketplaces such as TeachersPayTeachers.com. While the former flavor of tweets could be welcomed by fellow educators, the latter would likely not be appreciated by some users.

Implications for Future Research

Research to date has paid scant attention to spam's presence and effects on online educator communities, and there is fertile ground for additional work in this area. Here, our practical metrics and suggestions for their use have focused on volume-related spam common to extreme forms of participation inequality in online social groups. The primary implication of this paper for researchers is a call to attend to the impact that spam can have on their datasets and subsequent analysis. By employing the techniques we have suggested to better isolate "the signal" from "the noise," researchers can zero-in on the effects of spam, make more accurate comparisons, and add confidence to their findings.

Beyond these methodological implications, there remain numerous avenues for future study. Researchers could study whether spam impacts differ depending on user demographics such as race, gender, age, and years of teaching experience. Future studies could explore spam's effects at the level of the space or community. For instance, do certain hashtags thrive despite the existence of spam, while other hashtag spaces or communities experience significant damage? Research could also investigate the strategies educators employ to evaluate users and information in spam-heavy Twitter contexts. Additionally, given the commonplace use of multiple social media platforms, research that analyzes similarities and differences in spam across platforms (e.g., Twitter, Facebook and Pinterest) could prove beneficial. Spam likely functions in some different ways across different technologies. For instance, Twitter is largely unmoderated and unstructured compared to a closed Facebook group or an online discussion board like those hosted on Reddit (Staudt Willet and Carpenter 2019), where at least some spam may be weeded out by moderators.

Implications for Practice

Our findings may have fewer direct implications for practice than for research, but there are valuable practical benefits to the increased attention to spam that we have argued for here. Given that many scholars, inservice teachers and preservice teachers (PSTs) have received encouragement from some quarters to engage in professional social media activities, preparation for spam is needed. For example, female scholars would likely benefit from awareness of the online harassment that other scholars have experienced and the ways that those scholars have responded to such harassment (Veletsianos et al. 2018).

Because some teacher education programs require their PSTs to use Twitter (e.g., Carpenter and Morrison 2018; Cook and Bissonnette 2016; Luo and Clifton 2017; Luo et al. 2017), teacher educators and instructional designers also may benefit from being aware of how spam affects educational spaces and, subsequently, educators' experiences. It is not difficult to imagine how encountering excessive spam may obscure the potential benefits of professional networking and collaboration via Twitter. It may therefore behoove teacher educators to make the teachers with whom they work aware of strategies for identifying and managing the presence of different kinds of Twitter spam. Given the diversity of education-focused Twitter hashtags that exist, teacher educators may also consider recommending certain hashtags that feature less spam overall, or less of certain types of spam. For example, #Edchat seems to contain a high proportion of self-promotional tweets from self-stylized educational gurus, and very few replies (Staudt Willet 2019), which may limit the hashtag's utility to PSTs and novice teachers seeking personalized advice or informal mentoring.

Inservice teachers can also benefit from awareness and consideration of the various types of spam in education-focused Twitter hashtag spaces. Different hashtags appear to attract varied quantities and qualities of spam, and teachers may therefore want to assess the relative signal-to-noise ratio in the hashtags relevant to them. With some hashtags, the content that teachers post may be less likely to be seen or responded to given the volume of competing spam content. Because hashtags are user-created, educators and educator communities may therefore also consider instances when it could be beneficial to develop new hashtags that would potentially avoid the attention of spammers—at least for a period of time. The existence of many flavors of spam also means that educators who use Twitter for their own learning and/or for classroom applications with their students must develop the skills needed to evaluate the information and messages encountered on Twitter. For instance, educators should consider the ways that commercial motivations could be more subtly reflected in some Twitter posts, such as through connections between tweeters and educational technology products that can only be discovered by careful attention to an account's Twitter profile or visiting the website to which a profile links (Staudt Willet 2019).

Conclusion

We suggest that addressing spam should be a critical step for research studying online educator spaces and communities. Researchers should acknowledge—and potentially explore—the kinds of impacts spam may have on their data and describe steps they take to mitigate those impacts, even when spam is not the central concern of their research. This would advance understanding of the roles of social media in education by allowing for better comparisons of results across studies. Given the large amounts of educator professional activities that occur online, spam's impacts are important for educational researchers, teacher educators, and instructional designers to consider. Spam likely affects educators' social media use and it additionally creates challenges for researchers as they collect and analyze data. In many cases, spam can confound understanding of the phenomenon of interest. In this paper, we have described practical metrics and a holistic decision-making strategy that should prove useful and accessible for many educational technology researchers. We assert the value of human decision-making regarding spam and suggest that scholars should keep in mind the diversity of spam and the various ways that educators may analyze and react to it.

Compliance with Ethical Standards

Disclosure of Potential Conflicts of Interest The authors declare that they have no conflict of interest.

Research Involving Human Participants and/or Animals This article contains no studies with animals performed by the authors. All actions performed in studies involving human participants were conducted in accordance with the ethical standards of the relevant institutional review board, and also with the 1964 Helsinki declaration and its later amendments or comparable ethical standards.

Informed Consent This paper conducts secondary analysis of data collected during two earlier studies (Carpenter et al. 2018; Staudt Willet 2019). Those studies, as well as this current one, included only publicly available data from the social media platform Twitter that were collected unobtrusively. The data are described only in aggregate; we do not point to individual users in any identifiable way and generally focus on behavior trends among a corpus of several thousand participants on social media. As such, informed consent was deemed unnecessary by the respective Institutional Review Boards.

References

boyd, d. (2014). *It's complicated: The social lives of networked teens*. New Haven: Yale University Press.

Britt, V. G., & Paulus, T. (2016). "Beyond the four walls of my building": A case study of #Edchat as a community of practice. *American Journal of Distance Education, 30*, 48–59. https://doi.org/10.1080/08923647.2016.1119609 .

Brunton, F. (2013). *Spam: A shadow history of the internet*. Cambridge: MIT Press.

Burgess, D. (2012). *Teach like a pirate*. San Diego: Dave Burgess Consulting.

Carpenter, J. (2015). Preservice teachers' microblogging: Professional development via twitter. *Contemporary Issues in Technology and Teacher Education, 15*(2), 209–234.

Carpenter, J. P., & Harvey, S. (2019). "There's no referee on social media": Challenges in educator professional social media use. *Teaching and Teacher Education, 86*, 102904.

Carpenter, J. P., & Krutka, D. G. (2014). How and why educators use twitter: A survey of the field. *Journal of Research on Technology in Education, 46*(4), 414–434. https://doi.org/10.1080/15391523.2014.925701.

Carpenter, J. P., & Morrison, S. A. (2018). Enhancing teacher education… with twitter? *Phi Delta Kappan, 100*(1), 25–28.

Carpenter, J., Tani, T., Morrison, S. & Keane, J. (2018). Exploring the education twitter Hashtag landscape. In E. Langran & J. Borup (Eds.), proceedings of Society for Information Technology & teacher education (SITE) international conference (pp. 2230-2235). Washington, D.C., United States: Association for the Advancement of computing in education (AACE).

Carpenter, J.P., Koehler, M.J., Staudt Willet, K.B., & Greenhalgh, S.P. (2019). Spam, spam, spam, spam: Methodological considerations and challenges for studying educators' twitter use. In K. Graziano (Ed.), proceedings of Society for Information Technology & teacher education (SITE) international conference 2019 (pp. 2702-2711). Las Vegas, NV: Association for the Advancement of computing in education (AACE). Retrieved from https://www.learntechlib.org/p/208033/.

Chen, C., Zhang, J., Xiang, Y., Zhou, W., & Oliver, J. (2016). Spammers are becoming "smarter" on Twitter. *IT Professional, 18*(2), 66–70. https://doi.org/10.1109/MITP.2016.36 .

Cook, M. P., & Bissonnette, J. D. (2016). Developing preservice teachers' positionalities in 140 characters or less: Examining microblogging as dialogic space. *Contemporary Issues in Technology and Teacher Education, 16*(2), 82–109.

Dede, C. (2006). *Online professional development for teachers: Emerging models*. Cambridge: Harvard Education Press.

Donath, J. (2007). Signals in social supernets. *Journal of Computer-Mediated Communication, 13*, 231–251. https://doi.org/10.1111/j.1083-6101.2007.00394.

Evans, P. (2015). Open online spaces of professional learning: Context, personalisation and facilitation. *TechTrends, 59*(1), 31–36.

Gao, F., & Li, L. (2017). Examining a one-hour synchronous chat in a microblogging-based professional development community. *British Journal of Educational Technology, 48*, 332–347.

Goodyear, V. A., Parker, M., & Casey, A. (2019). Social media and teacher professional learning communities. *Physical Education and Sport Pedagogy, 24*(5), 421–433. https://doi.org/10.1080/17408989.2019.1617263.

Greenhalgh, S. P. (2018). *Spaces and their social frontiers: Using community dimensions to distinguish between teacher-focused hashtags on twitter (doctoral dissertation)*. East Lansing: Michigan State University.

Hsieh, B. (2017). Making and missing connections: Exploring Twitter chats as a learning tool in a preservice teacher education course. *Contemporary Issues in Technology and Teacher Education, 17*(4), 549–568.

Karpf, D. (2012). Social science research methods in internet time. *Information, Communication & Society, 15*, 639–661. https://doi.org/10.1080/1369118X.2012.665468.

Kimmons, R., & Veletsianos, G. (2018). Public internet data mining methods in instructional design, educational technology, and online learning research. *TechTrends, 62*(5), 492–500. https://doi.org/10.1007/s11528-018-0307-4.

Kraut, R. E., & Resnick, P. (2011). *Building successful online communities: Evidence-based social design*. Cambridge: MIT Press.

Kwak, H., Lee, C., Park, H., & Moon, S. (2010). What is Twitter, a social network or a news media?. In *Proceedings of the 19th International Conference on World Wide Web (WWW 2010)* (pp. 591-600). New York: Association for Computing Machinery (ACM).

Lin, P.-C., & Huang, P.-M. (2013). A study of effective features for detecting long-surviving Twitter spam accounts. In *The 15th international conference on advanced communications technology (ICACT), technical proceedings 2013* (pp. 841–846). Piscataway: IEEE.

Luo, T., & Clifton, L. (2017). Examining collaborative knowledge construction in microblogging-based learning environments. *Journal of Information Technology Education: Research, 16*, 365–390.

Luo, T., Sickel, J., & Cheng, L. (2017). Preservice teachers' participation and perceptions of Twitter live chats as personal learning networks. *TechTrends, 61*, 225–235. https://doi.org/10.1007/s11528-016-0137-1 .

Macià, M., & García, I. (2016). Informal online communities and networks as a source of teacher professional development: a review. *Teaching and Teacher Education, 55*, 291–307. https://doi.org/10.1016/j.tate.2016.01.02.

Mowbray, M. (2014). Automated Twitter accounts. In K. Weller, A. Bruns, J. Burgess, M. Mahrt, & C. Puschmann (Eds.), *Twitter and society* (pp. 183–194). New York: Peter Lang.

Nagle, J. (2018). Twitter, cyber-violence, and the need for a critical social media literacy in teacher education: A review of the literature. *Teaching and Teacher Education, 76*, 86–94. https://doi.org/10.1016/j.tate.2018.08.014 .

Prestridge, S., Tondeur, J., & Ottenbreit-Leftwich, A. T. (2019). Insights from ICT-expert teachers about the design of educational practice: The learning opportunities of social media. *Technology, Pedagogy and Education, 28*, 157–172. https://doi.org/10.1080/1475939X.2019.1578685 .

Rosenberg, J. M., Greenhalgh, S. P., Koehler, M. J., Hamilton, E., & Akcaoglu, M. (2016). An investigation of state educational twitter Hashtags (SETHs) as affinity spaces. *E-Learning and Digital Media, 13*, 24–44. https://doi.org/10.1177/2042753016672351.

Rosenberg, J. M., Greenhalgh, S. P., Wolf, L. G., & Koehler, M. J. (2017). Strategies, use, and impact of social media for supporting teacher community within professional development: The case of one urban STEM program. *Journal of Computers in Mathematics and Science Teaching, 36*, 255–267.

Shelton, C., & Archambault, L. (2018). Discovering how teachers build virtual relationships and develop as professionals through online teacherpreneurship. *Journal of Interactive Learning Research, 29*, 579–602.

Staudt Willet, K. B. (2019). Revisiting how and why educators use twitter: Tweet types and purposes in #Edchat. *Journal of Research on Technology in Education, 51*(3), 273–289. https://doi.org/10.1080/15391523.2019.1611507.

Staudt Willet, K.B., & Carpenter, J.P. (2019). Educators on the front page of the internet: Education-related subreddits as learning spaces. In K. Graziano (Ed.), proceedings of Society for Information Technology & teacher education (SITE) international conference 2019 (pp. 2787-2795). Las Vegas, NV: Association for the Advancement of computing in education (AACE). Retrieved from https://www.learntechlib.org/p/208040/.

Tour, E. (2017). Teachers' self-initiated professional learning through personal learning networks. *Technology, Pedagogy and Education, 26*, 179–192. https://doi.org/10.1080/1475939X.2016.1196236.

Tucker, L. (2019). Educational professionals' decision making for professional growth using a case of Twitter adoption. *TechTrends, 63*(2), 133–148.

Twitter (n.d.). Report spam on Twitter. Retrieved from https://help.twitter.com/en/safety-and-security/report-spam

Veletsianos, G. (2017). Three cases of hashtags used as learning and professional development environments. *TechTrends, 61*(3), 284–292.

Veletsianos, G., Houlden, S., Hodson, J., & Gosse, C. (2018). Women scholars' experiences with online harassment and abuse: self-protection, resistance, acceptance, and self-blame. *New Media & Society, 20*, 4689–4708. https://doi.org/10.1177/1461444818781324.

Wang, B., Zubiaga, A., Liakata, M., & Procter, R. (2015). Making the most of tweet-inherent features for social spam detection on Twitter. *CEUR Workshop Proceedings*, *1395*, 10–16. Retrieved from https://arxiv.org/abs/1503.07405

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.